

Appellants' Brief on Appeal
S/N 10/671,935
Docket: YOR920030330US1 (YOR.485)

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BOARD OF PATENT APPEALS AND INTERFERENCES**

In re Application of

Gustavson, et al.

Serial No.: 10/671, 935 Group Art Unit: 2193

Filed: September 29, 2003 Examiner: Do, Chat C.

For: METHOD AND STRUCTURE FOR PRODUCING HIGH PERFORMANCE
LINEAR ALGEBRA ROUTINES USING A SELECTABLE ONE OF SIX POSSIBLE
LEVEL 3 L1 KERNEL ROUTINES

Commissioner of Patents
Alexandria, VA 22313-1450

APPELLANTS' BRIEF ON APPEAL

Sir:

Appellants respectfully appeal the rejection of claims 1-20 in the Office Action mailed on October 4, 2007. A Notice of Appeal was timely filed on January 3, 2008.

I. REAL PARTY IN INTEREST

The real party in interest is International Business Machines Corporation, assignee of 100% interest of the above-referenced patent application.

II. RELATED APPEALS AND INTERFERENCES

There are no other appeals or interferences known to Appellants, Appellants' legal representative or Assignee which would directly affect or be directly affected by or have a bearing on the Board's decision in this appeal.

III. STATUS OF CLAIMS

Claims 1-20 stand rejected under 35 U.S.C. § 101 as allegedly directed to non-statutory subject matter. Claims 1, 6, 7, 12, 13, and 18 stand rejected under 35 U.S.C. § 102(b) as allegedly anticipated by co-inventor Gustavson's own prior publication ("Superscalar GEMM-based Level 3 BLAS – The On-going Evolution of a Portable and High-Performance Library"). Claims 3, 4, 9, 10, 15, and 16 stand rejected under 35 U.S.C. § 103(a) as allegedly unpatentable over Gustavson, further in view of US Patent 6,357,041 to Pingali et al. Claims 19 and 20 stand rejected under 35 U.S.C. § 103(a) as allegedly unpatentable over Gustavson, further in view of "PLAPACK: Parallel Linear Algebra Package Design Overview" by Philip (Alpatov) et al.

Claims 1, 5-7, 11-13, 17, and 18 stand rejected under nonstatutory obviousness-type double patenting over claims 21 and 22 of co-pending application 10/671,934. Claims 3, 4, 9, 10, 15, and 16 stand rejected under nonstatutory obviousness-type double patenting over claims 21 and 22 of co-pending application 10/671,934, further in view of Pingali.

All the above rejections are being appealed.

IV. STATUS OF AMENDMENTS

A Request for Reconsideration Under 37 CFR §1.116 was filed on December 4, 2007. In the Advisory Action mailed on December 19, 2007, the Examiner indicated that the arguments in the Request for Reconsideration Under 37 CFR §1.116 were not persuasive and that the rejections of record were maintained for all claims.

V. SUMMARY OF CLAIMED SUBJECT MATTER

As explained at lines 7-11 of page 4 of the specification, the conventional wisdom for linear algebra processing considers that only one kernel type is available for matrix multiplication. However, as explained at lines 11-14 of page 5, such limitation of having

a single kernel available for matrix multiplication forces data copying that limits efficiency of the multiplication processing.

The claimed invention, on the other hand, provides a method to reduce and/or eliminate such data copying by allowing a selection of an optimal kernel for the processing, as selected based on which matrix would most optimally reside in L1 cache.

Bases in the specification for the claims:

1. (Rejected) A method of improving at least one of speed and efficiency when executing a level 3 dense linear algebra processing on a computer (lines 5-7 of page 3; lines 5-7 of page 4), said method comprising:

automatically setting an optimal machine state on said computer for said processing by selecting an optimal matrix subroutine from among a plurality of matrix subroutines stored in a memory that could alternatively perform a level 3 matrix multiplication processing (lines 9-11 of page 3; lines 5-10 of page 4; lines 12 of page 10).

7. (Rejected) An apparatus (200, Fig. 2), comprising:

a memory (221, Fig. 2) to store matrix data to be used for a processing in a level 3 dense linear algebra program;

a processor to perform said processing (211, Fig. 2; line 18 of page 13 through line 16 of page 14); and

a selector (211, Fig. 2) to select an optimal one of a plurality of possible matrix subroutines to that could alternatively perform said processing, thereby automatically setting said apparatus into an optimal machine state to perform said processing (lines 9-11 of page 3; lines 5-10 of page 4; lines 12 of page 10).

13. (Rejected) A machine-readable storage medium (500, Fig. 5) tangibly embodying a program of machine-readable instructions executable by a digital processing apparatus to

perform a method of improving at least one of speed and efficiency when executing a linear algebra subroutine on a computer, said method comprising:

selecting an optimal matrix subroutine from among a plurality of matrix subroutines that can alternatively perform a level 3 matrix multiplication processing, thereby automatically setting said computer into an optimal machine state for performing said level 3 matrix multiplication processing (lines 9-11 of page 3; lines 5-10 of page 4; lines 12 of page 10).

19. (Rejected) A method of providing a service involving at least one of solving and applying a scientific/engineering problem, said method comprising at least one of (line 17 of page 26 through line 12 of page 27):

using a linear algebra software package that improves at least one of speed and efficiency to performs one or more matrix processing operations, wherein said linear algebra software package achieves the improved speed or efficiency by selecting an optimal matrix subroutine from among a plurality of matrix subroutines that alternatively can perform a matrix multiplication processing, thereby automatically setting a computer into an optimal machine state for performing said matrix multiplication processing;

providing a consultation for solving a scientific/engineering problem using said linear algebra software package;

transmitting a result of said linear algebra software package on at least one of a network, a signal-bearing medium containing machine-readable data representing said result, and a printed version representing said result; and

receiving a result of said linear algebra software package on at least one of a network, a signal-bearing medium containing machine-readable data representing said result, and a printed version representing said result.

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

Appellant presents the single following ground for review by the Board of Patent Appeals and Interferences:

GROUND 1: The Non-statutory Subject Matter Rejection under 35 USC §101 for Claims 1-20;

GROUND 2: The Anticipation Rejection for Claims 1, 6, 7, 12, 13, and 18, based on Co-inventor Gustavson's Previous Publication "Superscalar GEMM-based Level 3 BLAS – The On-going Evolution of a Portable and High-Performance Library";

GROUND 3: The Obviousness Rejection for Claims 3, 4, 9, 10, 15, and 16, based on Co-inventor Gustavson's Previous Publication "Superscalar GEMM-based ...", further in view of US Patent 6,357,041 to Pingali et al.;

GROUND 4: The Obviousness Rejection for Claims 19 and 20, based on Co-inventor Gustavson's Previous Publication "Superscalar GEMM-based ...", further in view of Philip Alpatov et al.; and

GROUND 5: The Double Patenting Obviousness Rejections for Claims 1, 5-7, 11-13, 17, and 18, based on Claims 21 and 22 of Copending Application 10/671,934, and for Claims 3, 4, 9, 10, 15, and 16, based on these two claims of Copending Application 10/671,934, further in view of Pingali.

VII. ARGUMENTS

GROUND 1: The Non-statutory Subject Matter Rejection under 35 USC §101 for Claims 1-20;

The Examiner's Position

The Examiner alleges that the claimed invention is directed to non-statutory subject matter. In the Advisory Action mailed on December 19, 2007, the Examiner alleges:

“The examiner respectfully submits that the claims do not clearly or inherently disclose the increasing speed and efficiency in terms of the hardware performance, but rather the increasing speed and efficiency is in term[s] of mathematical operation or computations. Thus [the] claims are directed to non-statutory subject matter. Further, the claims appear to preempt every substantial practical application of the idea embodied by the claim and there is no cited limitation in the claims that breathes sufficient life and meaning into the preamble so as the limit it to a particular practical application rather than being so broad and sweeping as to cover every substantial practical application of the idea embodied therein. Finally, the specification clearly discloses in page 24 that the machine-readable media can be [definitely] non-tangible medium as signal-bearing media as a whole which is clearly and definitely non-statutory.

Appellants' Position

Appellants begin by briefly and specifically addressing each of the above-recited points in order.

First, relative to the Examiner's contention that hardware performance is not being addressed in the claim language, Appellants respectfully submit that the claim language of even independent claim 1 clearly refers to “setting an optimal machine state on said computer”, thereby clearly establishing a connection to a tangible machine. It is well established case law that a computer programmed for a specific task constitutes a unique machine.

Second, Appellants note that “breathing sufficient life and meaning into the preamble” is an issue of patentable weight of the claim preamble language, not statutory

subject matter. Moreover, the claim limitations are not required to describe specific practical applications, as alleged by the Examiner.

Third, preemption is the entire purpose of a patent claim and of the US patent system. However, Appellants note that the present invention makes no attempt to keep the public from using the older, less efficient processing methods.

Fourth, relative to the Examiner's final sentence recited above, the Examiner points to no case holding to support this statement. From Appellants' perspective, the closest case law would appear to be *In re Nuijten*, 500 F. 3d 1346 (Fed. Cir. 2007), and, as explained in more detail below, the facts of that case are clearly distinguished from those of the present claimed invention.

Turning now to more generally addressing this statutory subject matter rejection, in paragraph 13. a., beginning on page 10 of the Office Action, the Examiner argues that "... the claims do not explicitly disclose a practical application of the optimal subroutine to perform matrix multiplication. Basically, the claims just disclose a method of selecting a subroutine from a set of subroutine[s] to perform a matrix multiplication. The improvement of speed/efficiency would not constitute as concrete, useful, and tangible as required under 35 U.S.C. 101."

In response, Appellants respectfully disagree and submit that the placing of the machine into an optimal state by selecting one of possible alternative kernels to perform a given processing inherently provides the advantage over conventional methods (wherein only one kernel is available, regardless of its optimality) of increasing speed and efficiency. Appellants further submit that increasing processing speed and efficiency is exactly the type of results one would desire from a patent and this result is even expressly mentioned in the independent claims.

Therefore, Appellants simply disagree with the Examiner's position that the present invention fails to satisfy the "useful, concrete and tangible result" standard of review for statutory subject matter, if this test is considered appropriate to apply to the method claims.

Moreover, it also noted that the claims include apparatus and Beauregard-type claims, and these claims are clearly addressed to statutory subject matter, even if the method claims were to be ultimately deemed as directed to non-statutory subject matter.

In paragraph 13. b. beginning on page 11 of the Office Action, the Examiner argues "... that the claim language recites a machine-readable storage medium tangibly embodying the program but not as a tangible machine readable storage medium embodying the program as alleged by the appellant. Further, the specification page 24 lines 10-15 does suggest that the machine readable medium can be non-tangible medium such as digital and analog and communication links and wireless. Clearly, the machine readable storage medium claims are directed to non-tangible medium."

In response, Appellants respectfully submit that it is the Examiner who summarily declares that the description at lines 10-15 of page 24 is both incorporated into the claim language and is non-statutory. Appellants respectfully disagree that the Examiner is necessarily correct.

First, it is brought to the Examiner's attention that the claim language itself limits the claim to "a machine-readable storage medium tangibly embodying a program of machine-readable instructions executable by a digital processing apparatus...." As such, this language clearly defines a "process" and is statutory by reason of being one of the four categories specifically itemized in 35 USC §101.

Second, the wording "machine-readable storage medium" clearly includes ROM and RAM containing the machine-readable instruction, as well as the standalone disks or diskettes of the Beauregard-type claims. Therefore, Appellants submit that, since the language clearly covers at least some statutory subject matter, the claimed invention is statutory. That is, the evaluation for statutory subject matter is the invention as a whole, not whether an Examiner is able to interpret the language as possibly including definitions for which case law arguably provides no clear answer.

Moreover, to the extent that the Examiner considers that this language includes the description on page 24 of the specification making reference to transmission media, Appellants respectfully point out that there is no case law in this regard whether a series of

Docket YOR920030330US1

machine-readable instructions that define a process to be executable by a machine is excluded from being statutory subject matter, particularly considering that, as mentioned above, a “process” is expressly identified as one of the four categories listed in 35 USC §101. To the extent that the Examiner considers that a transmission can “tangibly embody a program of machine readable instructions executable by a digital processing apparatus to perform a method of ...”, Appellants again point out that this transmission clearly defines a “process.”

The closest case law that would seem relevant to the Examiner’s concern would appear to be *In re Nuijten*, 500 F. 3d 1346 (Fed. Cir. 2007). However, in contrast to the present disputed claims, that case involved a claim specifically addressed to a “signal”, *per se*. Moreover, the claimed invention in *Nuijten* involved a watermark embedded in that signal, which the Court considered analogous to a product-by-process. In contrast, the claimed invention of the present application are clearly directed to a process, one of the four categories specifically identified in 35 USC §101.

Therefore, the Examiner’s position is clearly and improperly based upon an assumption that, if a “signal” is involved in any manner in defining a process, then such process is categorically eliminated from 35 USC §101. Appellants respectfully submit that neither Congress nor the Courts have so ruled and the Examiner fails to provide any support for this clear defiance of 35 USC §101.

Along this line, it is noted that neither Appellants nor the USPTO knows what mechanism might be utilized in infringing upon a process lawfully accorded patent protection. More particularly, with the advent of such technology as Bluetooth, it is easy to imagine a first machine/device controlling a second machine via transmitted signal that define and execute the claimed process steps.

It is further brought to the attention of the USPTO that all computers are already controlled by internal signals that are based on the electromagnetic spectrum, including those signals used to define the process steps being executed on that machine. Therefore, contrary to the Examiner’s implication, current machine-implemented process execution already relies upon “signals”, and Appellants respectfully submit that it is irrelevant

Appellants' Brief on Appeal
S/N: 10/671,935

whether those signals are external or internal to any specific machine, since it is the process that is being protected by these disputed claims.

Appellants hereby reaffirm on the record, in this Appeal Brief, that they are not waiving any right to protect this process in the future simply because the alleged infringer is able to point to a “signal” and be able to allege that the USPTO has declared that such signal thereby effectively shields the alleged infringer’s acts simply because such signal has rendered the process as nonstatutory, by declaration of the USPTO, and, therefore, incapable of being protected under 35 USC §101. Appellants do not believe that either Congress nor the Courts have given such blanket authority to the USPTO and that the USPTO should be extremely careful about enabling possible future infringement mechanism under cover of 35 USC §101 without having at least a reasonable fact pattern upon which to base its conclusion on non-statutory subject matter.

For the reasons stated above, the claimed invention is fully patentable over the reference, and the Board is respectfully requested to reconsider and withdraw this rejection.

GROUND 2: The Anticipation Rejection for Claims 1, 6, 7, 12, 13, and 18, based on Co-inventor Gustavson's Previous Publication "Superscalar GEMM-based Level 3 BLAS – The On-going Evolution of a Portable and High-Performance Library"

The Examiner alleges that co-inventor Gustavson's prior publication "Superscalar GEMM-based Level 3 BLAS – The On-going Evolution of a Portable and High-Performance Library" teaches the claimed invention defined by claims 1, 6, 7, 12, 13, and 18, and, when combined with the teachings of Pingali, renders obvious claims 3, 4, 9, 10, 15, and 16, and when combined with the teachings of Philip, renders obvious claims 19 and 20.

Appellants submit, however, that co-inventor Gustavson, as one of the authors of the cited primary reference, has declared unequivocally, in his previous response, that this publication described ways to write other level 3 BLAS in terms of DGEMM and featured only a single kernel and the use of data copying.

In contrast, the present invention describes the potential use of any of six kernel routines (one of which can be selected as optimal, particularly in view of one or more of others of the techniques described in the remaining six co-pending applications) and newer forms of data copying called "register blocking" (see co-pending Application S/N 10/671,888, corresponding to Attorney Docket YOR920030169US1). There is no suggestion in the Gustavson publication of using a selected one of six possible kernels, and the Examiner fails to point out specific locations reasonably demonstrating such plurality of selectable kernel subroutines.

That is, the Examiner points to section 3 in pages 208-209 and section 3.2 in pages 210-211 and the first four lines under the introduction section on page 207.

In response, Appellants respectfully submit that none of these locations even suggest the availability of alternative kernels, let alone selecting an optimal kernel from among six possible kernels.

That is, section 3 on pages 208-209 states:

“3 Superscalar GEMM-based level 3 BLAS

To approach peak performance on state-of-the-art superscalar microprocessors it is necessary to attain extensive register reuse. In general, multiple calls to the level 1 and level 2 BLAS routines prohibit an efficient register reuse.

Recently, Kågström and Ling announced the first version of the superscalar GEMM-based level 3 BLAS. They have also developed a superscalar DGEMM that currently is used with the library. The superscalar library has essentially the same overall structure, with similar blocking, as the regular GEMM-based level 3 BLAS. The main difference in the design is that all calls to underlying level 1 and level 2 BLAS have been removed. As before, the dominating part of all floating point operations take place in calls to DGEMM. The remaining computations that take care of triangular diagonal blocks are handled by "in-line" code optimized for efficient register reuse.”

Section 3.2 on pages 210-211 states:

“3.2 Improved performance for the superscalar library

In the current release of the superscalar GEMM-based level 3 BLAS, 4 x 4 unrolling is used for the C matrix in DGEMM and 4 x 2 unrolling is used in the remaining routines. As for the GEMM-based model implementations all references are stride one which is implemented using work arrays and data copying prearranged so that the DGEMM kernel will run close to peak performance. The extra data copying allows the superscalar library to handle so called "critical" leading dimensions as well [9, 10]. The Fortran source code is publically available from netlib, see www.netlib.org/blas/gemm_based/ssgemmbased.tgz’.

Performance results from the GEMM-based level 3 BLAS performance benchmark on an IBM PowerPC 604 processor (112 MHz, IBM SP, SMP node) show substantial improvements for the current release of the superscalar library:

Docket YOR920030330US1

DSYMM DSYRK DSYR2K DTRMM DTRSM

+3% +28% +2% +23% +25%

These percentage numbers are for square matrices of size 500 x 500. We obtain up to 80% improvement for small matrices (32 x 32). The improvements are mainly for the routines that called level 2 routines in the model implementations [9, 10]. The GEMM-based algorithms for DSYMM and DSYR2K do not call any level 2 routines. The calculations are transformed to level 3 GEMM operations by copying the symmetric subblocks stored in triangular format to general full format subblocks in work arrays [11].

The ATLAS [12] and PHiPAC projects [3] use the superscalar GEMM-based level 3 BLAS together with their own automatically tuned DGEMM to provide a complete set of level 3 BLAS in double precision. The ATLAS project reports impressive performance results for several different machines where the combination of the superscalar GEMM-based level 3 BLAS and ATLAS DGEMM is often faster than the vendor supplied level 3 BLAS, see 'www.netlib.org/atlas'."

The first four lines in the introduction page 207 state:

"1 Introduction

The level 3 Basic Linear Algebra Subprograms (BLAS) [4] are a de facto standard for various matrix multiply and triangular system solving computations and are successfully used as building blocks for the development of high-performance dense linear algebra library software."

Nowhere in the above-recited passages is there even a suggestion of alternative kernels or a selection of an optimal kernel from among a plurality of kernels that could alternately be used, and the Examiner is respectfully requested to point out specific lines intended to support his position. Co-inventor and co-author Gustavson states emphatically that this paper had no suggestion whatsoever of such alternative kernel selection.

Seemingly in response, in the Advisory Action mailed on December 19, 2007, the Examiner points to the above-recited “*DSYMM DSYRK DSYR2K DTRMM DTRSM +3% +28% +2% +23% +25%*” as supporting his position. However, these subroutines are not described in this section as being viable alternates for all matrix data or even for each other, thereby failing to satisfy the plain meaning of the claim language. The benefit recited in this section is clearly described as relative to subroutines in the existing library, not relative to each other, as the Examiner seems to imply, and does not suggest that any of these subroutines is an optimal alternative subroutine to the others, let alone by the method further articulated in various dependent claims.

Therefore, Appellants respectfully submit that the rejection currently of record fails to establish a *prima facie* rejection for either anticipation or obviousness, since it is fundamentally flawed by failing to provide a key element of the independent claims.

The Examiner relies upon secondary reference Pingali and tertiary reference Philip Alpatov for reasons unrelated to overcoming this basic deficiency of the primary reference, so that neither of these references overcome the deficiency of the primary reference.

Hence, turning to the clear language of the claims, in the Gustavson publication there is no teaching or suggestion of: “...automatically setting an optimal machine state on said computer for said processing by selecting an optimal matrix subroutine from among a plurality of matrix subroutines stored in a memory that could alternatively perform a level 3 matrix multiplication processing”, as required by independent claim 1. The remaining independent claims have similar language.

Therefore, Appellants submit that there are elements of the claimed invention that are not taught or suggested by Gustavson's prior publication, and the Board is respectfully requested to reconsider and withdraw this rejection.

Appellants' Brief on Appeal
S/N: 10/671,935

GROUND 3: The Obviousness Rejection for Claims 3, 4, 9, 10, 15, and 16, based on Co-inventor Gustavson's Previous Publication "Superscalar GEMM-based ...", further in view of US Patent 6,357,041 to Pingali et al.; and

Relative to the rejections based on combining secondary reference Pingali with the Gustavson publication, Appellants submit that this secondary reference fails to overcome the deficiency of the primary reference, and the Examiner does not allege otherwise, so that all of claims 3, 4, 9, 10, 15, and 16, are also clearly patentable over this publication, even if combined with this secondary references.

GROUND 4: The Obviousness Rejection for Claims 19 and 20, based on Co-inventor Gustavson's Previous Publication "Superscalar GEMM-based ...", further in view of Philip Alpatov et al.

Relative to the rejection for claim 19 and 20, based on combining Pingali and Philip Alpatov with the Gustavson publication, Appellants also submit that neither secondary reference Pingali nor tertiary reference Philip Alpatov overcomes the deficiency of the primary reference, and the Examiner does not allege otherwise, so that both claims 19 and 20 are also clearly patentable over this publication, even if combined with these two references.

GROUND 5: The Double Patenting Obviousness Rejections for Claims 1, 5-7, 11-13, 17, and 18, based on Claims 21 and 22 of Copending Application 10/671,934, and for Claims 3, 4, 9, 10, 15, and 16, based on these two claims of Copending Application 10/671,934, further in view of Pingali.

Claims 1, 5-7, 11-13, 17, and 18 stand rejected under nonstatutory obviousness-type double patenting over claims 21 and 22 of co-pending application S/N 10/671,934, and claims 3, 4, 9, 10, 15, and 16 stand rejected under nonstatutory obviousness-type double patenting over these claims 21 and 22 of co-pending application S/N 10/671,934, further in view of US Patent 6,357,041 to Pingali et al.

In response, Appellants again respectfully submit that co-pending application S/N 10/671,934 relates to a specific technique of streaming of data for level 3 matrix multiplication processing, not to the selection of an optimal subroutine for performing the processing. These procedures are clearly patentably distinct by reason of providing two distinctly different results, as evidenced by the different independent claims in the two applications.

That is, claims 21 and 22 of co-pending application S/N 10/671,934 respectively depend off of an independent claim that requires a determination of which matrix will reside in which cache layer. These two dependent claims 21 and 22 have to be interpreted as the combination of first determining which matrices reside on the various cache levels followed by the step of selecting two kernels from six possible kernels to perform a level three matrix multiplication processing.

In contrast, the independent claims of the present application S/N 10/671,935 define the entirely different and unrelated process of determining which one of alternative kernel to use as the subroutine for the matrix processing. None of the rejected claims in the present application '935 addresses the determination of a second kernel for the processing, as required by the independent claims of co-pending application '934.

Along this line, it is noted that the rejection currently of record fails to reasonably demonstrate any suggestion to determine even one optimal kernel, let alone two optimal

Docket YOR920030330US1

Appellants' Brief on Appeal
S/N: 10/671,935

kernels. Therefore, Appellants submit that it could hardly be considered obvious to select a second optimal kernel (as required by dependent claims 21 and 22 of '934), if it has not even been demonstrated to select a first optimal kernel.

The double patenting rejections currently of record provide no objective evidence or rationale to support this conclusion of obviousness, contrary to the requirement of the recent US Supreme Court holding in *KSR*: "There must be some articulated reasoning with some rational underpinning to support the legal conclusion of obviousness", *KSR Int'l v. Teleflex, Inc.*, 127 S. Ct. 1727, 1741, 82 USPQ2d 1385, 1396 (2007). The double patenting rejections consist of conclusory statements only; there is no analysis or rationale in these rejections, as required by the *KSR* holding.

Therefore, Appellants respectfully submit that these rejections for double patenting fail to meet the initial burden of a *prima facie* obviousness rejection, and the Examiner is respectfully requested to reconsider and withdraw these rejections.

IX. CONCLUSION

In view of the foregoing, Appellants submit that claims 1-20, all the claims presently pending in the application, are clearly enabled and patentably distinct from the prior art of record and in condition for allowance. Thus, the Board is respectfully requested to remove all rejections of claims 1-20.

Please charge any deficiencies and/or credit any overpayments necessary to enter this paper to Assignee's Deposit Account number 50-0510.

Respectfully submitted,



Dated: March 3, 2008

Frederick E. Cooperrider
Reg. No. 36,769

McGinn Property Law Group, PLLC
8231 Old Courthouse Road, Suite 200
Vienna, VA 22182-3817
(703) 761-4100
Customer Number: 21254
Docket YOR920030330US1

CLAIMS APPENDIX

The claims, as reflected upon entry of the Amendment Under 37 CFR §1.111 filed on July 16, 2007, are shown below:

1. (Rejected) A method of improving at least one of speed and efficiency when executing a level 3 dense linear algebra processing on a computer, said method comprising:

automatically setting an optimal machine state on said computer for said processing by selecting an optimal matrix subroutine from among a plurality of matrix subroutines stored in a memory that could alternatively perform a level 3 matrix multiplication processing.

2. (Rejected) The method of claim 1, wherein said computer includes an L1 cache, said method further comprising:

determining a size of each of matrices involved in said matrix multiplication; and selecting one of said matrices to reside in an L1 cache, based on said determined size,

wherein said selecting a matrix subroutine comprises determining which of said matrix subroutines is consistent with said matrix selected to reside in said L1 cache .

3. (Rejected) The method of claim 1, wherein said matrix subroutine comprises a substitute of a subroutine from LAPACK (Linear Algebra PACKage).

4. (Rejected) The method of claim 3, wherein said substitute LAPACK subroutine comprises a Basic Linear Algebra Subroutine (BLAS) Level 3 L1 cache kernel.

5. (Rejected) The method of claim 1, wherein said selecting a matrix subroutine comprises an aspect of a generalized matrix streaming process in which matrix data is stored in multiple levels of computer memory, including a matrix block stored in an L1 cache and matrix data of two other matrices stored in at least one higher level of cache, such that said matrix data of said two other matrices is systematically streamed into said matrix multiplication processing through said L1 cache.

6. (Rejected) The method of claim 1, wherein said plurality of matrix subroutines comprises six possible matrix subroutines that could alternatively be used for said level 3 matrix multiplication processing.

7. (Rejected) An apparatus, comprising:
 - a memory to store matrix data to be used for a processing in a level 3 dense linear algebra program;
 - a processor to perform said processing; and
 - a selector to select an optimal one of a plurality of possible matrix subroutines to that could alternatively perform said processing, thereby automatically setting said apparatus into an optimal machine state to perform said processing.

8. (Rejected) The apparatus of claim 7, further comprising an L1 cache, wherein said selector makes the selection by:

determining a size of each of matrices involved in said level 3 processing; and
selecting one of said matrices to reside in said L1 cache, based on said determined sizes,

wherein said selecting a matrix subroutine comprises determining which of said matrix subroutines is consistent with said matrix selected to reside in said L1 cache.

9. (Rejected) The apparatus of claim 7, wherein said matrix subroutine comprises a substitute of a subroutine from LAPACK (Linear Algebra PACKage).

10. (Rejected) The apparatus of claim 9, wherein said substitute LAPACK subroutine comprises a Basic Linear Algebra Subroutine (BLAS) Level 3 L1 cache kernel.

11. (Rejected) The apparatus of claim 7, wherein said selector for selecting a matrix subroutine includes a storage for storing matrix data in multiple levels of computer memory and a mechanism for streaming said matrix data into said matrix multiplication process.

12. (Rejected) The apparatus of claim 7, wherein said plurality of matrix subroutines comprises six possible matrix subroutine kernel types.

13. (Rejected) A machine-readable storage medium tangibly embodying a program of machine-readable instructions executable by a digital processing apparatus to perform a method of improving at least one of speed and efficiency when executing a linear algebra subroutine on a computer, said method comprising:

selecting an optimal matrix subroutine from among a plurality of matrix subroutines that can alternatively perform a level 3 matrix multiplication processing, thereby automatically setting said computer into an optimal machine state for performing said level 3 matrix multiplication processing.

14. (Rejected) The machine-readable storage medium of claim 13, wherein said digital processing apparatus includes an L1 cache, said method further comprising:

determining a size of each of matrices involved in said matrix multiplication processing; and

selecting one of said matrices to reside in an L1 cache, based on said determined size,

wherein said selecting a matrix subroutine comprises determining which of said matrix subroutines is consistent with said matrix selected to reside in said L1 cache.

15. (Rejected) The machine-readable storage medium of claim 13, wherein said matrix subroutine comprises a substitute for a subroutine from LAPACK (Linear Algebra PACKage).

16. (Rejected) The machine-readable storage medium of claim 15, wherein said substitute LAPACK subroutine comprises a Basic Linear Algebra Subroutine (BLAS) Level 3 L1 cache kernel.

17. (Rejected) The machine-readable storage medium of claim 13, wherein said selecting a matrix subroutine comprises an aspect of a generalized matrix streaming process in which matrix data is stored in multiple levels of computer memory, including a matrix block stored in an L1 cache and matrix data of two other matrices stored in at least one higher level of cache or other memory, such that said matrix data of said two other matrices is systematically streamed into said matrix multiplication processing through said L1 cache.

18. (Rejected) The machine-readable storage medium of claim 13, wherein said plurality of matrix subroutines comprises six possible kernel type subroutines.

19. (Rejected) A method of providing a service involving at least one of solving and applying a scientific/engineering problem, said method comprising at least one of:

using a linear algebra software package that improves at least one of speed and efficiency to performs one or more matrix processing operations, wherein said linear algebra software package achieves the improved speed or efficiency by selecting an optimal matrix subroutine from among a plurality of matrix subroutines that alternatively

can perform a matrix multiplication processing, thereby automatically setting a computer into an optimal machine state for performing said matrix multiplication processing;

providing a consultation for solving a scientific/engineering problem using said linear algebra software package;

transmitting a result of said linear algebra software package on at least one of a network, a signal-bearing medium containing machine-readable data representing said result, and a printed version representing said result; and

receiving a result of said linear algebra software package on at least one of a network, a signal-bearing medium containing machine-readable data representing said result, and a printed version representing said result.

20. (Rejected) The method of claim 19, wherein said matrix subroutine comprises a Basic Linear Algebra Subroutine (BLAS) Level 3 L1 cache kernel from LAPACK (Linear Algebra PACKage).

Appellants' Brief on Appeal
S/N: 10/671,935

EVIDENCE APPENDIX

None

RELATED PROCEEDINGS APPENDIX

None